Cyclic Twill-Woven Objects

Ergun Akleman[1], Jianer Chen[2], YenLin Chen[2], Qing Xing[3], Jonathan L. Gross[4]

[1]Visualization Department, Texas A&M University
[2]Computer Science Department, Texas A&M University
[3]Architecture Department, Texas A&M University
[4]Computer Science Department, Columbia University

Abstract

Classical (or biaxial) twill is a textile weave in which the weft threads pass over and under two or more warp threads, with an offset between adjacent weft threads to give an appearance of diagonal lines. This paper introduces a theoretical framework for constructing *twill-woven objects*, i.e., cyclic twill-weavings on arbitrary surfaces, and it provides methods to convert polygonal meshes into twill-woven objects. It also develops a general technique to obtain exact triaxial-woven objects from an arbitrary polygonal mesh surface.

1. Motivation

Figure 1: Three biaxial twill-woven objects obtained by applying Catmull-Clark subdivision to the same initial mesh.

Beyond its use in fabric design, weaving provides a wide variety of ways to create surface patterns that can be embodied in sculpture and in innovative architectural design. We focus here on twill-weaving, which can provide strength, durability, and water-resistance, along with interesting diagonal patterns. We describe methods to convert polygonal meshes into twill-woven sculptures.

It has recently been shown (Akleman, et al. [3, 4]) how any given polygonal mesh can be transformed into objects woven from ribbons of varying width, such that the ribbons cover the underlying surface almost completely, except for small holes. The ribbons can be manufactured inexpensively by using laser-cutters on thin metal sheets. The corresponding plain-woven sculptures are constructed physically by weaving these metal ribbons. Mallos [26] has created large-scale plain-woven objects.

With the design and construction of more and more unusually shaped buildings, the computer graphics community has started to explore new methods to reduce the cost of the physical construction for large shapes. Most of the currently suggested methods focus on reduction of the number of differently shaped components [31, 11, 12]. There exists a contemporary interest among architects to explore weaving as an alternative construction method [23, 13, 18] based on traditional bamboo-woven housing [19, 20, 29]. This suggests how weaving with ribbons from thin metal sheets can also be useful for economical construction of complicated shapes.
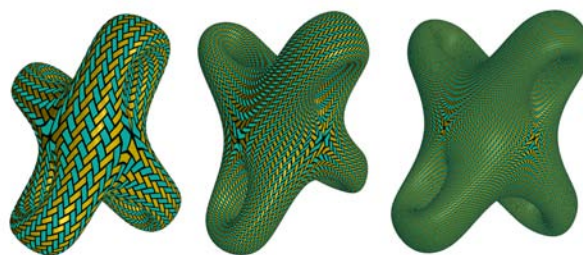
2. Introduction

It has recently been described [3] how any arbitrary twist of the edges of an extended graph rotation system induces a cyclic plain-weaving on the corresponding surface. The method works for all meshes and is very simple: just twisting all edges in the same way. On the other hand, in order to construct other weaving structures we had to use two different types of twisted edges, characterize conditions to design the desired weaving structures, and develop algorithms to create exact and approximate versions of the weaving structures. Herein we extend the mathematical model to twill-weaving, which is used in fabrics such as denim or gabardine. Classical twill is a biaxial textile weave in which each weft (filling) thread passes over and under two consecutive warp threads, and each row is obtained from the row above it by a shift of one unit to the right or to the left. The shift operation creates a diagonal pattern that adds visual appeal to a twill fabric or to a twill-woven object, as illustrated in Figure 1.

In what follows, we define a twill-weaving as a type

of cyclic weaving structure on general surfaces. Based on this definition, we identify three mesh conditions that are collectively necessary and sufficient to obtain a twill-weaving from a given mesh. In fact, many meshes do not satisfy these three conditions, which implies that it is impossible to obtain an exact twill for them. Intuitively, we may expect that a mostly-$(4, 4)$ mesh, i.e., a mesh with large areas of quadrilaterals with 4-valent vertices, would admit a reasonably good twill. Indeed, we have developed an edge-coloring algorithm that will create an exact twill whenever the mesh is twillable. Even if the mesh is not twillable, the output of the algorithm satisfies most of the twill conditions, and it exhibits the characteristic diagonal pattern for mostly-$(4, 4)$ meshes, as shown in Figure 11(c).

Our generalized definition of twill leads us to identify a previously unknown weaving pattern that we call triaxial twill. Triaxial twill patterns are created from meshes that are populated with $(3, 6)$ regions (i.e., triangles with 6-valent vertices). Such meshes can be obtained by triangular schemes such as mid-edge subdivision [25] or $\sqrt{3}$ subdivision [24]. We prove that every mesh obtained by mid-edge subdivision [25] is twillable. Triaxial twill patterns are visually interesting and reminiscent of some of the tilings of M. C. Escher, as shown in Figure 11(d). We note that a triaxial twill does not demonstrate the characteristic diagonal pattern of classical biaxial twill.

Obtaining a biaxial twill that exhibits diagonal patterns depends on having a proliferation of $(4, 4)$ regions in the mesh. Quad-remeshing schemes such as the Catmull-Clark [6] or Doo-Sabin subdivisions [10] can achieve that proliferation, and they can make the number of crossings in each cycle divisible by 4, an important property of twill weaving. In §8, we show the existence of meshes that continue to be biaxially twillable after application of quad-remeshing schemes.

3. Textile Twill

Most textile weaves are 2-way weaves, also called *biaxial* weaves. They consist of row and column strands, called *weft* and *warp* respectively, at right angles to each other. They are also *2-fold*, which means that there are never more than two strands crossing each other. The popularity of biaxial weaving comes from the fact that most textile weaving structures are manufactured using loom devices by interlacing the two sets of strands at right angles to each other. The basic purpose of any loom device is to hold the warp strands under tension, so that the weft strands can weave under and over warp strands to create a fabric. Using a loom, it is possible to manufacture a wide variety of weaves by raising and lowering different warp strands.

Grunbaum and Shephard [16] formally investigated the mathematics behind these 2-way, 2-fold woven fabrics. They coined the phrase *isonemal fabrics* [17] to describe 2-way, 2-fold fabrics that have a transitive symmetry group acting on the strands. Twill weaving belongs to a certain family of isonemal fabrics in which each weft row of length n (the *period*) is obtained from the weft row immediately above it by a cyclical shift of s units (the *offset*) to the right, for some fixed value of the parameter s, such that n and s are relatively prime. If $s = \pm 1$, the fabric is called twill [17]. More generally, the resulting fabrics are called (n, s)-fabrics. This family of fabrics also includes plain-weaves (with $n = 2$, $s = 1$) and satins.

Twills are widely used in clothing fabrics, for instance, in denim or gabardine. Their characteristic diagonal pattern that makes the weaving visually appealing. Since twill-weaving uses fewer crossings than plain-weaves, the yarns in twill-woven fabrics can move more freely than the yarns in plain-woven fabrics. This property makes twill-weaving softer, more pliable, and better draped than plain-weaving. Twill fabrics also recover better from wrinkles than plain-woven fabrics. Moreover, yarns in twill-weaving can be packed closer. This property makes the twill-woven fabrics more durable and water-resistant, which is a reason why twill-fabrics are often used for sturdy work clothing or for durable upholstery.

A special family of twills is characterized by two integers a and b, where a is the number of over-crossings, and b the number of under-crossings of a weft thread as it crosses the warp threads. Each twill-weaving pattern in this family can be expressed by a triple $a/b/s$, where $s = \pm 1$ and where $b/a/\text{-}1$ is a 90°-rotated version of $a/b/1$. For instance, $2/2/1$ and $2/2/\text{-}1$ in Figure 2 define 90°-rotated versions of exactly the same twill structures. Similarly, $3/1/1$ and $1/3/\text{-}1$ define exactly the same twill-weaving structures. This is expected since one over in weft must correspond to one under in warp and one under in weft must correspond to one over in warp. Thus, the total number of overs and the total number of unders should be equal.
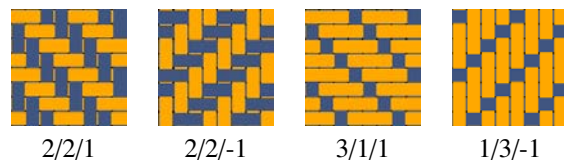

| 2/2/1 | 2/2/-1 | 3/1/1 | 1/3/-1 |

Figure 2: All possible biaxial twills for period $n = 4$.

In this paper, we focus mainly on $2/2/1$ twill, simply called $2/2$ twill, but our results generalize to other $a/b/1$ twills. The $2/2$ twill pattern is quite common, and its characteristic pattern of twill is not limited to fabrics. Brick walks and hardwood floor tiles sometimes exhibit twill patterns. (See Figure 3(c).)

This work can also be useful to study the covering of an arbitrary surface with hexagonal tiles. The visual relationship between the $2/2$ twill pattern and the regular hexagonal tiling is shown in Figure 4. [28] [27]
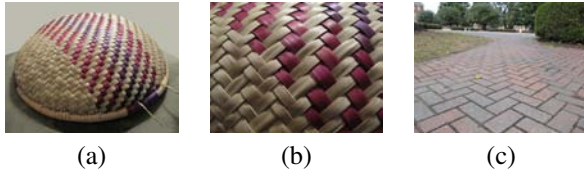
(a)   (b)   (c)

Figure 3: Non-fabric examples of twill pattern. (a) is a twill woven basket from Mozambique, and (b) shows its detail. (c) is an hexagonal tiling that exhibits the same diagonal pattern as twill weaving.
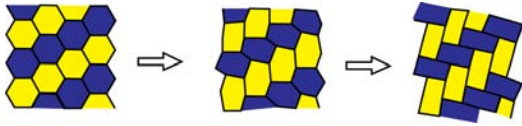

Figure 4: A regular hexagonal tiling and a related 2/2 twill weaving.

In this paper, we are concentrating on topological and aesthetic properties of twills, rather than exploring the physical aspects of twill-weaving. Our goal is to obtain twill patterns on surfaces and to understand the impact of the initial mesh structures. Unlike plain-weaving, which results in an isotropic checkerboard-like pattern, twill-weaving results in an anisotropic structure.

4. General Twill-Weaving on Surfaces

In this section, we generalize twill weaving from planar weaves to weaves on arbitrary surfaces. We define a *cyclic weaving* on an arbitrary orientable manifold surface S_o to be a projection of a link L to S_o, such that there are no triple intersections of the link image at a single point on S_o [3]. Under this definition, cyclic weaves on general surfaces are 2-fold but not necessarily 2-way weaves. That is, a cyclic weave need not have a clear distinction of warp and weft cycles. In some exceptional cases, like those shown in Figure 1, the characteristic diagonal patterns of twill-weaves, as shown in Figure 2, can be obtained by coloring warp and weft threads with different colors. On the other hand, when the warp and weft distinction does not happen, it is still possible to obtain an approximate twill, which exhibits a strong diagonal pattern in most places (see Figure 11).

Similar to the terminology of knot theorists (for knot theory, see [9, 22]), we call the part of a cycle between any two consecutive crossings a *segment*, and we call the two crossing points the *ends* of the segment. Two segments are called *adjacent* if they share a crossing. A *gap* is a "hole" in the weave, something like a region of a graph drawing that allows edge-crossings. For instance, in Figure 5(b), the region r of the surface that is partially bounded by the sequence of segments a, b, c is a gap.
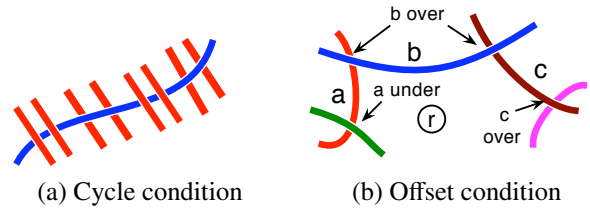

(a) Cycle condition   (b) Offset condition

Figure 5: 2/2-twill weaving conditions.

Definition 4.1. General 2/2-twill weaving. *A cyclic weave on a surface S_o is a 2/2-twill if it satisfies the following conditions:*

- Cycle condition: *In a complete traversal of every cycle, one must alternatingly encounter pairs of over-crossings and pairs of under-crossings. See Figure 5(a).*

- Offset condition: *Let a, b, and c be three consecutive segments in a traversal of the boundary of a gap, such that both ends of segment b are of the same type (i.e, either both over or both under). Then the other ends of segments a and c (i.e., the ends not touching segment b) must be of different types from each other. See Figure 5(b).*

Akleman, et al. [2] introduces the concept of extended graph rotation systems, which facilitates a practical description of a cyclic weave as a mapping of a link L on the orientable surface S_o, and develops a projection algorithm to convert the weaving cycles to 3D threads, such as ribbons or yarns. In the next section, we describe how cyclic weaves are specified by extended graph rotation systems and how they are constructed by means of the projection algorithm.

5. Extended Graph Rotation Systems

Let G be a graph (without any self-loops). A *rotation* at a vertex v is a cyclic ordering of the edges incident at v. A *rotation system* for the graph G is a collection of rotations, one for each vertex of G. By the Heffter-Edmonds Theorem (see [14]), there is a one-to-one correspondence between graph rotation systems and graph embeddings on orientable surfaces, in which the rotation at a vertex corresponds to the ordering of the incident edges embedded around the vertex on the surface.

In an extended graph rotation system (EGRS), an edge is viewed as a thin rectangular strip that can be twisted clockwise or counter-clockwise, in helical sense (see Figure 6). The sides of the strips provide two "strands" in the corresponding weaving pattern. For an "untwisted edge", these strands are "parallel" to the mesh-edge, and for a "twisted edge", they both cross over the mesh-edge and one over the other. The type of the crossover depends on the helical twist sense (i.e., clockwise or counter-clockwise). If an arbitrary subset of edges of a mesh on an orientable surface S_o is

twisted in the same helical sense, then the EGRS induces a cyclic plain-weaving on S_o [2]. We now show how controlling the twist-type of every edge makes it possible to construct different weaves, beyond simple plain-weaving .
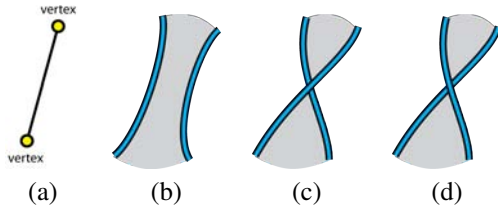

Figure 6: (a) An edge, as usually represented. (b) A visual interpretation of an untwisted EGRS edge. (c) A counter-clockwise twisted EGRS edge. (d) A clockwise twisted EGRS edge.

5.1. Specifying Weaves on Surfaces with EGRS

A rotation system $\rho_o(G)$ of a graph $G = (V, E)$ with no twisted edges uniquely determines a set W of closed walks that serve as the set of boundaries of regions of an embedding of that graph on an orientable surface S_o of Euler characteristic $|V| - |E| + |W|$. The set W is constructed by applying the *face-tracing algorithm*, and each closed walk in W is regarded as the boundary of a polygon. The surface S_o is obtained by pasting the polygons together. (See [14].) If we imagine that the graph G is in 3-space, then the polygons (if the polygons are sufficiently flexible) can be pasted to G so as to form the surface S_o in 3-space. We imagine further that the graph G is thickened to a "regular neighborhood" N in which the edges become the thin rectangular strips. Thus, we have a model in 3-space of a thickened graph on a closed surface.

Figure 7(a)(b)(c) illustrates this for the example of the octahedral graph on the sphere. Observe that the neighborhood N of the graph is a surface-with-boundary that lies on the sphere. Notice also that the components of $bd(N)$ form a link in 3-space, in which the components are completely unlinked.

Next let A be an arbitrary subset of edges of G. Choosing either the clockwise or counterclockwise sense and then twisting all the edges of A in that direction gives an *EGRS* $\rho(G)$. (Beyond the usual case of a single half-twist of an edge, an EGRS permits any number of half-twists, each specifying a unique weave.) Twisting the corresponding strips in the surface N has the effect of linking components of $bd(N)$. An effect of twisting the strips is to link some of the components of $bd(N)$ together, thereby forming a link L in 3-space, with a natural projection to the surface S_o, as shown in Figure 7(d)(e)(f). Applying the face-tracing algorithm to the EGRS $\rho(G)$ constructs the components of the link L. The link L mapped on S_o makes a cyclic weaving W
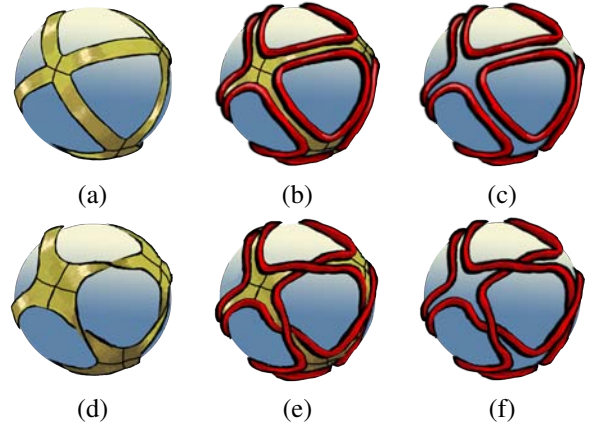

Figure 7: An example of creating a cyclic weaving on a surface with face-tracing. The mesh in (a) is an octahedron with untwisted edges, embedded in a sphere. As shown in (c), the eight face-boundary walks of the octahedron are initially unlinked in 3D. However, if some edges are twisted, as shown in (d), then some of the cycles represented by the boundary walks become merged together, and the resulting cycles are linked, as in (f).

on the surface S_o. We say that the cyclic weaving W is *induced* by the extended graph rotation system $\rho(G)$.

An important property of cyclic weaving, easily derived [3], is that every cycle c in a cyclic weaving induced by an extended graph rotation system has an even number of crossings. If c is a self-crossing, then every sub-cycle of c separated by a self-crossing has an even number of crossings (see Figure 8). This property holds, regardless of how we twist the edges, and it will be used in the rest of the paper.
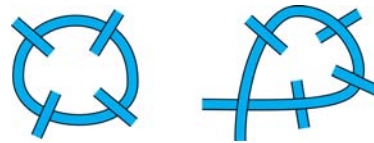

Figure 8: Every cycle and sub-cycle in a cyclic weaving induced by a EGRS has an even number of crossings.

5.2. General 2/2 Twill Conditions for Meshes

We now describe how the theory of extended graph rotation systems can be used to simplify our initial problem of specifying generalized twill-weaving on surfaces to a problem in edge-bicoloring, in such a way that the induced cyclic weaving resembles a 2/2 twill-weave with a strong diagonal pattern in most places. A local perspective on the edge-bicoloring is represented in Figure 9:

- blue means the edge is counter-clockwise twisted;

- red means the edge is clockwise twisted.

We will call a mesh *twillable* if there exists at least one edge-labeling solution such that the corresponding EGRS $\rho(G)$ induces a general 2/2 twill-weave.

Definition 5.1. *We now redefine a 2/2 twill-weaving in terms of an edge-bicoloring. Let ρ be the EGRS obtained by twisting all the edges of a mesh M_o.*

- Edge condition: *Any three edges consecutive in a "face boundary" induced by ρ must have two consecutive like-colored edges and the other edge colored differently. See Figure 9(a).*

- Face condition: *If two consecutive edges in a "face boundary" F_o in M_o are colored differently in ρ, then their neighboring two edges in F_o must be like-colored in ρ. See Figure 9(b).*

- Vertex condition: *If two edges consecutive in the rotation at a vertex v in M_o are colored differently in ρ, then their neighboring two edges in the rotation at v must be like-colored in ρ. See Figure 9(c).*
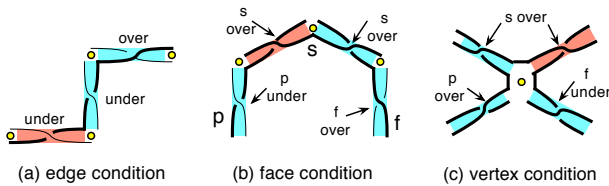

Figure 9: 2/2-twill mesh conditions.

Theorem 5.2. Twillable mesh conditions. *A mesh is twillable if and only if there exists a edge bicoloring that satisfies the edge, vertex, and face conditions given just above.*

Proof: As we described above, the two sides of an edge in the mesh correspond to two cycle segments in the corresponding weaving structure. In particular, a clockwise twist of the edge implies that the left segment under-crosses the right segment at their crossing point, while a counterclockwise twist of the edge implies that the left segment over-crosses the right segment at their crossing point (the sides of the segments are determined when we traverse the edge on the surface along a specific direction. Note that no matter which direction of the edge we are traversing, the above statements always hold true). Therefore, two consecutive segments in a cycle from two consecutive edges of the same twisting type in a face boundary induce two different types of crossings, while two consecutive segments in a cycle from two consecutive edges of different twisting types in a

face boundary induce two crossings of the same crossing type (note that a segment moves from one side to the other side after it passes through its crossing point).

Based on this observation, and comparing Figure 5 and Figure 9, we can easily prove the theorem now: the edge condition for meshes in definition 5.1 is just equivalent to the cycle condition for 2/2 twills in definition 4.1, while the face and vertex conditions for meshes in definition 5.1 are equivalent to the offset condition for 2/2 twills in definition 4.1. □

Figure 10 shows an example of a $(4, 4)$-grid (i.e quadrilaterals with 4-valent vertices) structure that is edge-bicolored to satisfy twillable mesh conditions. The edge bicoloring on the right defines a 2/2 twill-weave.
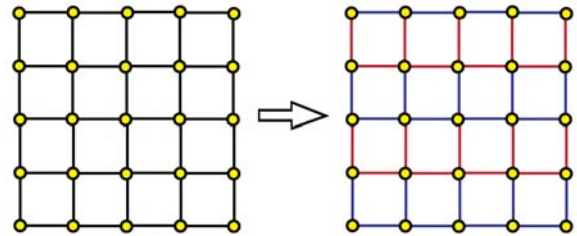

Figure 10: Applying twillable mesh conditions to a regular $(4, 4)$ mesh creates an anisotropic structure of twist types.

6. Edge Bicoloring with Voting

Each cycle in a twill must have the number of its crossings divisible by 4, because the pattern is 2-over, 2-under, and the numbers of over-crossings and under-crossings must be equal. Although many meshes do not admit an exact twill, for meshes with large sub-meshes of quadrilaterals with 4-valent vertices, it seems intuitively possible to obtain a reasonably good twill-weave in most places. Based on this intuition, we have developed and implemented an edge-bicoloring algorithm that can satisfy the twill conditions in most places. The resultant weavings can demonstrate strong diagonal patterns everywhere. This edge-bicoloring algorithm creates exact twill if the mesh is twillable. When we encounter an obstruction to the extension of a twill, we use a "voting" strategy that seeks to minimize the number of conflicts with the edge, face and vertex conditions.

6.1. Partitioning Threads into Bunches

In order to distinguish the cycles of the mesh from those of the induced weaving, we refer to the latter as *threads*. We observe that the mesh itself induces the threads. That is, since every edge of the mesh is to be twisted, the face-tracing algorithm yields the same set of threads, no matter which way any individual edge is colored or twisted.

A preliminary step before assigning colors (representing twist directions) to the edges is to partition the threads into bundles such that no cycle crosses another in the same bundle. (This is like vertex-coloring in graph theory.) For instance, if there exist only two bundles of threads, they are classified as warp and weft. The characteristic diagonal pattern of twill becomes visible when distinct colors are applied to warp and weft threads (see Figure 1). Even in the cases when partitioning into two bundles is not possible, giving all the threads in the same bundle the same color illuminates the structure of the weaving, as shown in Figure 11(c).

6.2. Breadth-first Extension with a Voting Strategy

Our edge bicoloring involves growing a sub-mesh of colored (i.e, twisted) edges by coloring *frontier edges*, i.e., edges that are incident on a vertex of the growing submesh, but are not yet colored. The color is chosen to meet all the twilling conditions, if possible. When exact twilling is not possible, we use a "voting" rule. Here is the procedure:

1. Initialize submesh S with an arbitrarily selected vertex v_0. Color all edges incident at v_0 (i.e., the first frontier set) in such a way that the vertex condition is satisfied (the edge and face conditions are vacuously satisfied), and add those edges to S, including their other endpoints.
2. Determine the set F of frontier edges for S.
3. Select a frontier edge from F for which one of the two possible colors satisfies all available edge, vertex and face conditions and assign that color. Continue such selections from F and color assignments as long as possible. The remaining edges are called "leftover edges".
4. Each leftover edge is potentially subject to four applications of the edge condition, since each of its two threads traverses a neighboring edge in each of two directions. It is similarly subject to four applications of the face condition and four applications of the vertex condition. We view each of these 12 conditions as having an equal "vote". When some neighboring edges are not yet colored, a vote may be unavailable. If all available conditions agree, the edge is colored accordingly. If there is any conflict, the decision is postponed.
5. Until every edge has been considered, return to Step 2.
6. Color a least conflicted (i.e., strongest majority) edge e by majority rule; then recalculate the votes for the neighboring edges. Continue choosing and coloring a least conflicted edge until all edges are colored.

6.3. Group Hierarchy in Voting

Partitioning the threads into bunches and assigning the same color to each thread in the same bunch helps to produce strong diagonal patterns. The voting power of bunch is proportional to the sum of the lengths of its threads, which permits a largest bunch greater power in defining the twill. This approach provides a consistent look as shown in Figure 11.

6.4. Weighted Voting

The vertex and face conditions are important to obtaining diagonal patterns. The edge condition guarantees only that the cycle alternates two-over and two-under. If the mesh is not twillable, it is helpful to give the vertex and face conditions higher weighting in the voting than the edge condition. Allowing 3/1 twill in few places helps to achieve strong diagonal patterns as it can be seen in Figure 11.

As mentioned earlier, if the mesh is twillable, then our edge-bicoloring algorithm constructs an exact twill. Fortunately, there exist a large number of twillable mesh families. We study mesh families that can provide exact biaxial twill in Section 8. In studying twillable meshes, we have discovered that there also exists another family of twill patterns, which we call triaxial twill. Interestingly, it is easy to obtain meshes that can produce exact triaxial twill, as explained in the next section.

7. Exact Triaxial Twill

Triaxial twill-weaves are created from meshes with $(3, 6)$ regions (i.e triangles with 6-valent vertices) predominating. Such meshes can simply be obtained (see Loop [25]) by *mid-edge subdivision*, which involves superimposing what graph theorists call a *medial graph* on an arbitrary mesh. As illustrated in Figure 12 (left), the procedure begins with a "parent mesh":

1. Insert a "new vertex" at the midpoint of each parent edge. This subdivides each parent edge into two "old edges".
2. Traverse the boundary of every face of the parent mesh, and in so doing, join each pair of consecutive new vertices by inserting a "new edge".
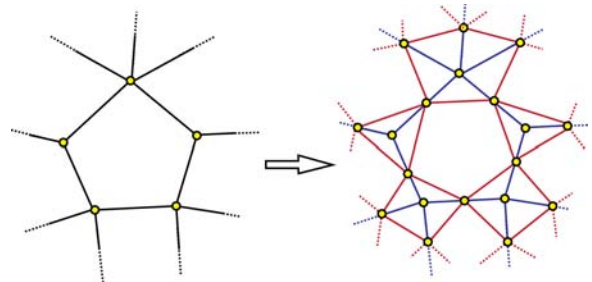3. Color every old edge blue and every new edge red (Figure 12 (right)) — or vice versa.

Figure 12: *Mid-edge subdivision always produces a twillable mesh.*

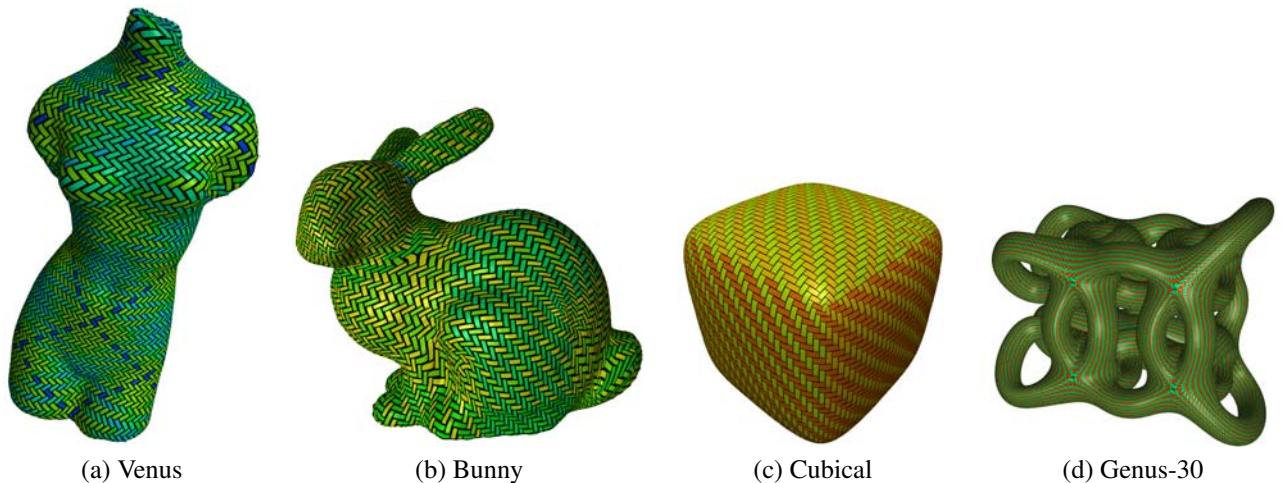(a) Venus   (b) Bunny   (c) Cubical   (d) Genus-30

Figure 11: Examples of Approximate biaxial cyclic twill woven objects constructed from polygonal meshes with the voting algorithm. Genus-30 example has its cycles partitioned into two bunches, even though the weaving does not satisfy twill conditions everywhere.

Theorem 7.1. *For any parent mesh on any orientable surface, the mid-edge subdivision scheme produces a twillable mesh.*

Proof: Let M be a mesh obtained by the mid-edge subdivision scheme on an arbitrary mesh. It suffices to prove that the coloring process on the mesh M given in step 3 above gives a bicoloring of the edges of M that satisfies the conditions of Theorem 5.2.

First observe the following: (1) if v is an old vertex in M, then all edges incident to v are old edges; (2) if v is a new vertex in M, then v is incident to exactly 6 edges, in which two are old edges, four are new edges, and in the rotation at v, the two old edges are separated by two consecutive new edges on each side; (3) each face in M either has all new edges on its boundary, or is a triangle made by one new edge and two old edges; and (4) no two faces whose boundaries are all new edges share a common edge.

Based on these observations, we can easily verify that the bicoloring given in step 3 satisfies the conditions of Theorem 5.2. For example, to see that the face condition is satisfied, let e_1 and e_2 be two consecutive edges in the boundary of a face F in M. If e_1 and e_2 are colored with different colors, then the face F must be a triangle whose boundary consists of one new edge and two old edges. Therefore, the other edge of F must be an old edge, which is neighboring to both e_1 and e_2 in the face boundary. Thus, the neighboring edges of e_1 and e_2 (which in this case are the same edge) must be colored with the same color. That is, the face condition in Theorem 5.2 is satisfied. The vertex condition and the edge condition can be verified similarly. □

Triaxial twill patterns are interesting, reminiscent of some M. C. Escher tilings (as in Figure 13). However, triaxial twill does not demonstrate the characteristic diagonal pattern.
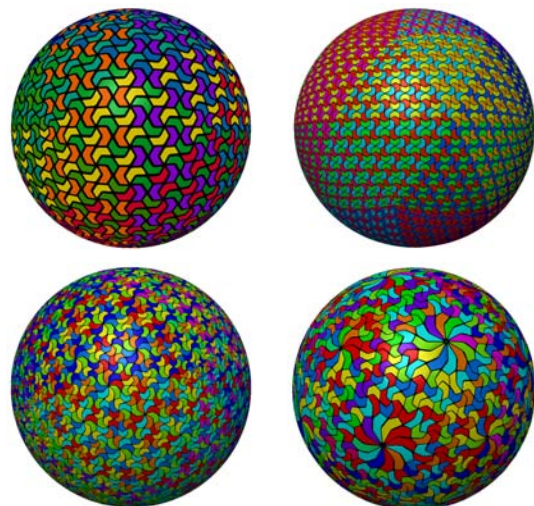

Figure 13: Examples of triaxial twill weaving on genus-0 surfaces.

8. Exact Biaxial Twill

To obtain the characteristic diagonal patterns, we use biaxial twill that is obtained from meshes populated by $(4, 4)$ regions (i.e., meshes with large areas of quadrilaterals with 4-valent vertices). The *simplest method* (Peters-Reif [30]) involves overlaying the medial graph. The *vertex-insertion method* (Catmull-Clark [6]) involves insertion of a new vertex at the center of a region, followed by a subdivision of each edge on the region boundary, and adjoining the resulting vertex to the vertex at the center of the region. The *corner-cutting method* (Doo-Sabin [10]) involves reapplication of the medial graph construction. These are collectively called *quad-remeshing schemes*. They can convert any mesh

into a quad mesh, either in which all faces are 4-sided, or in which all vertices are 4-valent. Moreover, iteratively applying such quad-remeshing schemes causes $(4, 4)$ regions to predominate.

If a quad mesh M_0 is twillable, then any mesh obtained from M_0 by these quad-remeshing schemes is also twillable. This hereditary twillability helps us to create an arbitrarily dense weaving, by iterative quad-remeshing.

Having threads whose numbers of crossings are multiples of 4 crossings is not sufficient to obtain a 2/2 twill. We must also create consistent offsets to obtain the twill look. To create offsets, twisting the edges so as to satisfy the vertex and face conditions in definition 5.1 is still important.

We have experimentally created a few such meshes and we have formally shown the existence of quad-meshes on every surface of genus larger than 0 in the form of $(4m, 4n, g)$[1] [1]. [Examples of twill-woven objects obtained by remeshing $(8, 4, 3)$ and $(16, 4, 4)$ meshes are shown in Figures 14 and 15, respectively. It is also possible to use semi-regular polyhedra to create twillable meshes. The twill-woven objects in Figure 16 come from descendant-twillable meshes that were produced by adding handles to $2n$-gonal bipyramids [33], such as an octahedron.

Figure 14: Diagonal twill-weaving patterns on a quad mesh obtained by remeshing a $(8, 4, 3)$ regular mesh.

Another important property of cyclic biaxial twill-woven objects is that two thread colors are sufficient, as shown in Figures 1, 15, 14 and 16. This is because the threads in an exact biaxial twill-weaving can be naturally partitioned into warp and weft thread bundles such that no two threads in the same bundle cross each other.

9. Implementation and Results

We have developed a system by implementing an edge bicoloring algorithm to obtain twill patterns. This algorithm provides exact twill-weaves for biaxial and triaxial twillable meshes and acceptable approximate

Figure 15: Diagonal twill-weaving patterns on a quad mesh obtained by remeshing a $(16, 4, 4)$ regular mesh.

twill weaves for other meshes. Our system can transform any given polygonal mesh into an approximate or exact cyclic twill-woven object. The system converts threads of the weave into ribbons or into yarns by using the projection algorithm described in [3]. Users can interactively change the thickness of ribbons or yarns. Figure 17 shows an example of an edge bicoloring and of the resulting cyclic woven objects, for thick and thin ribbons. All the twill woven-object images in this paper are direct screen captures from the system, and they were created in real-time. The colors of the threads are chosen based on the partitioning algorithm.

Bicolored edges   Thin ribbons   Thick ribbons

Figure 17: The correspondence between the bicolored edges and cyclic weave.

We use saturated colors to emphasize the twill structure. For choosing the colors, we employ two different strategies.

1. To get colorful yet balanced results as in Figure 1, we sample the hue space with equal angles. This strategy is useful to emphasize strong diagonal patterns, which are clearly visible in the case of non-self crossing cycles.
2. To get a subtle twill effect as in Figure 1(a), (b) and (c), we sample only a small portion of the hue space with equal angles. This is used in the case of self-crossing cycles. Despite the color similarities, the diagonal pattern remains evident in these images.

The projection algorithm always guarantees that the sizes of the ribbons are relative to the underlying polygons. Therefore, the widths of the ribbons vary in differ-

---

[1]The notation (n, m, g) denotes a genus-g regular mesh where all faces are n-sided and all vertices are m-valent [1, 7, 32].
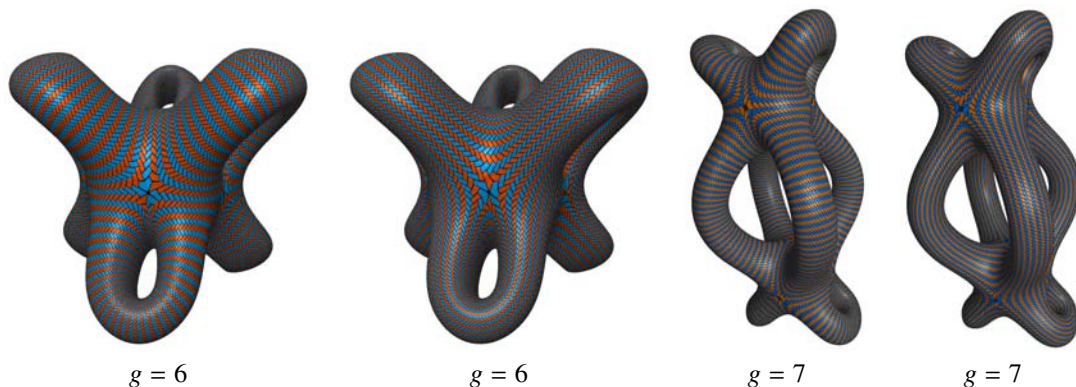
$g = 6$ $g = 6$ $g = 7$ $g = 7$

Figure 16: Two higher genus examples of biaxial twill woven objects that are induced by a quad mesh obtained from meshes produced by adding handles to $2n$-gonal bipyramids.

ent parts of the mesh. The projection method can create a sparse weaving, but, since twill becomes more visible in dense weaving, we use only dense woven objects in our examples. The projection method closes the gaps better with the mesh is obtained from a few iterations of a subdivision scheme. Since we already apply subdivision schemes to obtain denser meshes, the projection method covers the entire surface very smoothly, and the results look as if texture mapped onto smooth surfaces. Perhaps surprisingly, the results we have achieved cannot be obtained with texture mapping.

3D geometry allows us to achieve more realism in an interactive rendering. We can compute a real-time shadow, which can show the top ribbon's shadow on the bottom ribbon. We can also use a specularly reflected environment map, which looks significantly different in the top and bottom ribbons. Having 3D geometry also allows us to change the width of the ribbons in real time. Since we can easily unfold the ribbons, it is theoretically possible to construct these woven surfaces as physical shapes.

10. Conclusion and Future Work

We have introduced a theoretical framework for creating twill-woven objects based on extended graph rotation systems. We have formulated generalized twill conditions that apply to twills on any surface, not just on the plane. Based on the twill conditions, we have developed an edge bicoloring algorithm that produces acceptable approximate twill-weaves for any given mesh and exact twill for twillable meshes. Moreover, we have developed a procedure to convert an arbitrary mesh to generalized triaxial twill, which exhibits an interesting pattern that is significantly different from diagonal patterns. In our investigations so far, we have not located commercial textile looms for triaxial twill. However, triaxial twill patterns provide an interesting hexagonal tiling of surfaces.

We have shown, furthermore, that for any surface of positive genus, it is possible to construct meshes that

can be converted to a biaxial twill that exhibits diagonal patterns. Since the geometry is not important for twillability of a mesh, the same mesh structure can be used to cover any surface of the same genus. In addition, we have shown that these twill-weaving structures can be painted in two colors. We have identified the effect of subdivision schemes on the cyclic weaving structures.

Our results are not limited to twill. The relationship between mesh subdivision and weaving refinement can effectively be used for other weaving and knitting patterns. It may be useful to check some remeshing schemes that we have not analyzed, such as the remeshing scheme of $\sqrt{3}$ subdivision [24].

Our results also suggest that it is possible to provide consistent anisotropy over an arbitrary surface of positive genus, and it can be used to understand the limitations and requirements to create anisotropic structures on arbitrary surfaces. Another important implication of this result can be the characterization of the mesh structures that are parameterizable, since any consistent parametrization can be used to cover the surface with a consistent anisotropic pattern.

An interesting direction is to explore the links between plane symmetry patterns [15, 8] and weaving. There currently exists an interest in computer graphics for exploring symmetry patterns on surfaces such as N-way rotational symmetries (N-RoSy) fields [28, 27]. In particular, 4-RoSy fields have potential to design twill patterns on a given surface while moving singularities to invisible or natural positions.

We are currently researching on the material and constructing technicalities of weaving on an architectural scale. As we have discussed in section 1, plain or twill woven objects can be constructed from ribbons of varying width that are manufactured inexpensively by using laser-cutters on thin metal sheets. For future research, there are several practical issues that need to be resolved to make this happen. For instance, the unrolled flat ribbons are both long and wavy. Cutting such a ribbon as one continuous piece results in a huge amount of waste material. It is our experience that the physical construc-

tion of weaving is difficult with such long and wavy ribbons. Therefore we currently explore ways, such as shortening the strips into more manageable lengths, for easier construction of woven objects.

The venus and bunny meshes that are used to create approximate biaxial cyclic twill woven objects in Figure 11(a) and (b) are created by the Quadcover method [21], courtesy of Wenping Wang and Li Yupei. The rest of the models in the paper are created using TopMod3D [5].

References

[1] Akleman, E., Chen, J., 2006. Regular mesh construction algorithms using regular handles. In: Proc. IEEE Int. Conference on Shape Modeling and Applications, (SMI 06). pp. 171–181.

[2] Akleman, E., Chen, J., Gross, J., Xing, Q., 2009. Extended graph rotation systems as a model for cyclic weaving on orientable surfaces. Technical Report - Computer Science Department, Texas A&M University - TR09-203.

[3] Akleman, E., Chen, J., Xing, Q., Gross, J., 2009. Cyclic plain-weaving with extended graph rotation systems. ACM Transactions on Graphics; Proceedings of SIGGRAPH'2009, 78.1–78.8.

[4] Akleman, E., Chen, J., Xing, Q., Gross, J., 2010. Single-cycle plain-woven objects. Proceedings of IEEE International Conference On Shape Modeling And Applications (SMI) 2010, 90–99.

[5] Akleman, E., Srinivasan, V., Chen, J., Morris, D., Tett, S., 2008. Topmod3d: An interactive topological mesh modeler. In: Proc. of CGI'2008. pp. 10–18.

[6] Catmull, E., Clark, J., 1978. Recursively generated b-spline surfaces on arbitrary topological meshes. Computer Aided Design (10), 350–355.

[7] Conder, M., Dobcsnyi, P., 2001. Determination of all regular maps of small genus. Journal of Combinatorial Theory Series B 81 (2), 224–242.

[8] Conway, J. H., Burgiel, H., Goodman-Strauss, C., 2008. The Symmetries of Things. A. K. Peters.

[9] Crowell, R. H., Fox, R. H., 1977. Introduction to Knot Theory. Dover Books on Mathematics.

[10] Doo, D., Sabin, M., 1978. Behavior of recursive subdivision surfaces near extraordinary points. Computer Aided Design (10), 356–360.

[11] Eigensatz, M., Kilian, M., Schiftner, A., Mitra, N. J., Pottmann, H., Pauly, M., 2010. Paneling architectural freeform surfaces. ACM Transactions on Graphics (TOG) - Proceedings of ACM SIGGRAPH 2010 29 (4), 45.1–45.8.

[12] Fu, C.-W., Chi-Fu Lai, Y. H., Cohen-Or, D., 2010. K-set tilable surfaces. ACM Transactions on Graphics (TOG) - Proceedings of ACM SIGGRAPH 2010 29 (4), 44.1–44.8.

[13] Griffen, C. S., 2001. Weaving as an analogy for architectural design. ACSA Annual Meeting: http://faculty.philau.edu/griffenc/weaving_analogy.htm.

[14] Gross, J. L., Tucker, T. W., 1987. Topological Graph Theory. Wiley Interscience, New York.

[15] Grunbaum, B., Shephard, G., 1987. Tilings and Patterns. W. H. Freeman and Co, NY.

[16] Grunbaum, B., Shephard, G. C., 1980. Satins and twills: an introduction to the geometry of fabrics. Mathematics Magazine 53, 139–161.

[17] Grunbaum, B., Shephard, G. C., 1988. Isonemal fabrics. American Mathematical Monthly 95, 5–30.

[18] Housdan, J., 2009. Weaving architecture: An exploration of old and new materials and construction methods in washington d.c. Viriginia Tech, Thesis, http://scholar.lib.vt.edu/theses/available/etd-01152010-122829/unrestricted/Housdan_JJK_T_2009.pdf.

[19] Janssen, J., 1981. Bamboo in building structures. Ph. D. thesis. Eindhoven, (Holland): University of Technology.

[20] Janssen, J. J., 1995. Building with Bamboo: A Handbook. Intermediate Technology Publications, Publishing, Warwickshire, UK.

[21] Kalberer, F., Nieser, M., Polthier, K., 2007. Quadcover - surface parameterization using branched coverings 26 (3), 10–19.

[22] Kauffman, L. H., 1983. The Symmetries of Things. Dover Books on Mathematics.

[23] Kieran, S., Timberlake, J., 23 April 2003. Architectural weaving. Architecture Week 95, B1.1.

[24] Kobbelt, L., 2000. $\sqrt{3}$-subdivision. Proceedings of SIGGRAPH'2000, 103–112.

[25] Loop, C., 1987. Smooth subdivision surfaces based on triangles. Master's thesis, University of Utah.

[26] Mallos, J., 2009. How to weave a basket of arbitrary shape. Accepted to ISAMA'2009; also see http://jamesmallos.blogspot.com/.

[27] Nieser, M., Palacios, J., Polthier, K., Zhang, E., 2010. Hexagonal global parameterization of arbitrary surfaces. In: ACM SIGGRAPH ASIA 2010 Sketches. pp. 5:1–5:2.

[28] Palacios, J., Zhang, E., 2007. Rotational symmetry field design on surfaces. In: ACM SIGGRAPH, ACM Transactions on Graphics (TOG). pp. 55:1–55:10.

[29] Paudel, S., Lobovikov, M., 2003. Bamboo housing: market potential for low-income groups. Journal of Bamboo and Rattan 2 (4), 381–396.

[30] Peters, J., Reif, U., 1997. The simplest subdivision scheme for smoothing polyhedra. ACM Transactions on Graphics 16 (4), 420–431.

[31] Singh, M., Schaefer, S., 2010. Triangle surfaces with discrete equivalence classes. ACM Transactions on Graphics (TOG) - Proceedings of ACM SIGGRAPH 2010 29 (4), 46.1–46.8.

[32] Wijk, J. J. V., 2009. Symmetric tiling of closed surfaces: Visualization of regular maps. ACM Transactions on Graphics; Proceedings of SIGGRAPH'2009, 28.1–28.12.

[33] Williams, R., 1972. The Geometrical Foundation of Natural Structures. Dover Publications, Inc.